Chapter 1 / Capítulo 1

SAP Superintelligence Series
ISSN: 3144-5624
ISBN (complete work): 978-9915-9851-0-7

Superintelligence Series Volume 4: One Person Company: Artificial Intelligence, Autonomous Entrepreneurship and the Future of Business
ISBN (this volume): 978-9915-9966-3-9
DOI: 10.62486/978-9915-9966-3-9.Ch01
Pages: 1-11

Check for updates

© 2026 The authors. This is an open access book chapter distributed under the terms of the Creative Commons Attribution (CC BY) 4.0 License.

Business formation and the public release of generative artificial intelligence: an interrupted time-series analysis

La formación de empresas y la publicación pública de la inteligencia artificial generativa: un análisis de series temporales interrumpidas

Libardo Ramón Polanía[1]
Rómulo Esteban Montilla[2]

1 Universidad de la Amazonia. Florencia, Colombia.
2 St. Mary's University, Texas, Estados Unidos

ABSTRACT

Introduction: The public release of generative artificial intelligence in November 2022 promised to lower the cost of starting and running a firm, yet its effect on aggregate business formation remains unmeasured. Methods: We used monthly business applications from the US Census Bureau Business Formation Statistics, seasonally adjusted, from January 2018 through August 2026, and fitted an interrupted time-series model with a break at November 2022, estimating level and trend changes with Newey-West standard errors (12 lags) for the aggregate and for three knowledge-intensive sectors. Results: At the national level the immediate shift was −49,832 applications (p = 0.099) and the monthly trend slowed by 1,602 applications (p = 0.034), from a pre-release slope of 3,631. Sectoral responses diverged, with Information falling 2,123 applications at the break (p = 0.032) and Management of companies accelerating by 44 applications per month (p < 0.001). Conclusions: The release of generative artificial intelligence coincided with a slowdown in aggregate business formation and with divergent, sector-specific responses, a pattern consistent with skill-biased reallocation rather than a uniform startup surge.

Keywords: business formation; generative artificial intelligence; interrupted time series; entrepreneurship; United States

RESUMEN

Introducción: La publicación pública de la inteligencia artificial generativa en noviembre de 2022 prometía abaratar el costo de crear y operar una empresa, pero su efecto sobre la formación agregada de empresas sigue sin medirse. Métodos: Usamos las solicitudes mensuales de creación de empresas de la Business Formation Statistics del US Census Bureau, desestacionalizadas, de enero de 2018 a agosto de 2026, y ajustamos un modelo de series temporales interrumpidas con un corte en noviembre de 2022, estimando cambios de nivel y de tendencia con errores estándar de Newey-West (12 rezagos) para el agregado y para tres sectores intensivos en conocimiento. Resultados: A nivel nacional el cambio inmediato fue de −49 832 solicitudes (p = 0,099) y la tendencia mensual se desaceleró en 1 602 solicitudes (p = 0,034), desde una pendiente previa de 3 631. Las respuestas sectoriales divergieron, con Información cayendo 2 123 solicitudes en el corte (p = 0,032) y Gestión de empresas acelerándose en 44 solicitudes mensuales (p < 0,001). Conclusiones: La publicación de la inteligencia artificial generativa coincidió con una desaceleración de la formación agregada de empresas y con respuestas sectoriales divergentes, un patrón consistente con una reasignación sesgada por habilidades más que con un auge uniforme de nuevas empresas.

Palabras clave: formación de empresas; inteligencia artificial generativa; series temporales interrumpidas; emprendimiento; Estados Unidos

INTRODUCTION

The United States entered the 2020s with a business formation system that had already been losing dynamism for decades. Decker *et al.*[1] documented that the startup rate and the pace of employment reallocation declined steadily from the 1980s onward, and Akcigit and Ates[2] organized the stylized facts of that decline into a coherent account centered on falling knowledge diffusion and rising market power. The magnitude is large. Guzman and Stern[3] estimated that entrepreneurial quality, measured from founding conditions, deteriorated alongside quantity across 32 states between 1988 and 2014, and Sterk, Sedláček and Pugsley[4] showed that the nature of firm growth changed in ways that reduced the share of young firms in aggregate employment. The decline was not a one-off: Aguilera-Bravo, Casares and Khan[5] asked directly whether US business dynamism had recovered in the 2010s and concluded that the answer was at best partial, while Haltiwanger[6] traced the persistence of the slowdown into the twenty-first century. Business formation is the upstream margin through which new ideas become firms, and Decker *et al.*[7] showed that the responsiveness of entry to productivity shocks weakened precisely when dynamism fell.

The geography of this decline matters for what follows. Rubinton[8] linked the spatial pattern of business dynamism to skill-biased technical change, showing that the sectors and places most exposed to new technology reallocated entry rather than uniformly expanding it. This is the same mechanism that a general-purpose technology such as generative AI would be expected to trigger. Kim *et al.*[9] updated the US evidence with new data and confirmed that firm growth and entry have stagnated even as the economy has shifted toward intangible, knowledge-intensive activity, a shift that sets the stage for a technology whose main product is knowledge work itself. At the macro level, the decline has also been read as a symptom of economic complexity: Bueno and Ferraz[10] revisited creative destruction through that lens and found that the decline in entry is not uniform across the complexity of the activity, which anticipates the sectoral heterogeneity that a knowledge technology would amplify. The lesson drawn from this literature is that a shock to the cost of knowledge work should be expected to redistribute the rate of new-firm formation rather than to raise it uniformly across sectors.

The public release of generative artificial intelligence in November 2022 altered the expected cost structure of entrepreneurship. Large language models and image generators diffused within weeks to millions of users, and economists moved quickly to quantify their labor-market implications. Eloundou *et al.*[11] estimated that roughly 80 percent of the US workforce had at least 10 percent of their tasks exposed to large language models, while Noy and Zhang[12] found in a randomized field experiment that access to ChatGPT raised the productivity of professional writing tasks by about 40 percent and reduced time by a comparable margin. Brynjolfsson, Li and Raymond[13] reported similar gains among customer-support agents, with the largest benefits accruing to less experienced workers. Brown *et al.*[14] located these results within a broader theory of generative AI in business and management, arguing that the technology reconfigures how value is created rather than merely automating a fixed set of tasks, and Bauch *et al.*[15] organized the concrete use cases that firms actually deploy. These results imply that a prospective founder could produce a business plan, marketing copy, code and routine analysis at a fraction of the prior cost, which should, other things equal, lower the barrier to entry.

The entrepreneurship literature has begun to theorize this link but has not yet estimated it at the aggregate margin. Fossen, McLemore and Sorgner[16] reviewed the mechanisms through which artificial intelligence could either substitute for or complement entrepreneurs, and Dinlersoz, Dogan and Zolas[17] provided the first census-scale evidence that AI-using startups are more likely to grow, which raises the returns to entry. Mariani and Dwivedi[18] positioned generative AI as a general-purpose technology that shifts the innovation process itself, and the GAIN framework of[19] consolidated the emerging insights into a research agenda organized around opportunity recognition, resource assembly and scaling, while Li *et al.*[20] and Rajaram and Tinguely[21] showed that its adoption is associated with distinct entrepreneurial practices in small and medium enterprises. The institutional and category-level effects are equally salient: Grodal *et al.*[22] described how the generative AI category was socially constructed across organizational fields, Nzembayie and Urbano[23] theorized generative AI platforms as institutional catalysts of digital entrepreneurship, and Albashrawi[24] synthesized its role in decision-making across disciplines. What remains unresolved is the aggregate time-series question: whether the November 2022 release changed the level or the trend of business formation, and whether that change was uniform or concentrated in particular sectors.

The time-series question is consequential for both interpretations. If generative AI acts as

https://doi.org/10.62486/978-9915-9966-3-9.Ch01

a broad cost reduction, the release should produce a positive level shift and a steeper trend in national business formation. If it instead acts as a skill-biased shock, the aggregate response should be muted or negative while the sectoral responses diverge in sign. These two hypotheses are observationally distinguishable, and the distinction matters for policy: a broad surge would call for generic support, whereas a reallocation would call for support targeted at the sectors the technology displaces. The present analysis is designed to separate the two. This study estimates those quantities directly. The objective is to estimate the change in the level and trend of US business formation associated with the public release of generative AI, and to characterize sectoral heterogeneity across knowledge-intensive industries.

METHODS

DESIGN AND DATA SOURCE

This is an interrupted time-series analysis of monthly business formation applications, reported following the RECORD statement for research using routinely collected administrative data. The source is the Business Formation Statistics of the US Census Bureau, downloaded as the file bfs_monthly.csv from the Census Economic Programs without authentication. The series measures the number of Employer Identification Number applications filed each month that are likely to become employer businesses, and is published both seasonally adjusted and unadjusted; we used the seasonally adjusted national series (sa = A, geo = US). Monthly values were extracted from the wide-format release, in which each row is a sector, series and year and the twelve month columns hold the counts, and reshaped into a long monthly panel. We dropped the unadjusted series and retained only seasonally adjusted counts to remove within-year seasonality that would otherwise dominate the monthly variance. The break month was assigned to the post-release period because the release occurred within November 2022, so the level term captures the first month in which the technology was publicly available. No missing values were present across the 105-month window for any of the four series, and no imputation was therefore required.

SAMPLE AND PERIOD

We restricted the series to January 2018 through August 2026, the last month with complete data at the time of extraction, yielding 105 monthly observations. The start was chosen to isolate a stable pre-release baseline after the immediate post-2008 recovery and before the COVID-19 distortions of 2020; the model below absorbs residual pandemic shocks through the time and break terms. Observations before 2018 were excluded because the series extends to 2004 and the earlier decades embed a different dynamism regime documented by Decker *et al.*.(7)

VARIABLES

The outcome is the monthly count of business applications in series BA_BA, the broadest application measure. BA_BA counts all applications for an Employer Identification Number that are indicated to have a high likelihood of becoming employer businesses, and it is the widest of the Census application series; we preferred it to the high-propensity subset because a technology shock that lowers the cost of founding should show first in the broad margin of intent rather than in the narrower margin of near-certain employers. The exposure is time relative to the break, defined at November 2022, the month ChatGPT was released to the public and generative AI entered general use. We constructed a month index (t) counting from zero at January 2018, a post-release indicator equal to zero before November 2022 and one thereafter, and an interaction term equal to the number of months elapsed since the break, set to zero before the break. We analyzed the aggregate sector (TOTAL) and three knowledge-intensive sectors defined by the North American Industry Classification System: Information (NAICS 51), Professional, scientific, and technical services (NAICS 54), and Management of companies and enterprises (NAICS 55). These three sectors were selected because they concentrate the writing, programming, analysis and coordination tasks that the exposure literature identifies as most affected by generative AI, and because they differ in whether the technology substitutes for or complements the sector's core output.

https://doi.org/10.62486/978-9915-9966-3-9.Ch01

STATISTICAL ANALYSIS

We fitted, for each sector, an interrupted time-series model of the form

$$y = \beta 0 + \beta 1\, t + \beta 2\, post + \beta 3\, (t - t^*) \cdot post + \varepsilon$$

where β2 is the immediate level change at the break and β3 is the change in the monthly trend after the break. This segmented-regression specification is the standard approach for an exposure with a known onset, and Wang *et al.*[25] clarified the exact interpretation of its coefficients, which we follow: the level term captures the discontinuity at the break and the interaction term captures the change in slope. The specification is a member of the family reviewed by,[26] who distinguished the traditional segmented model from the alternative counterfactual approaches available for environmental and policy time series. Parameters were estimated by ordinary least squares with Newey-West heteroskedasticity- and autocorrelation-consistent standard errors using 12 lags, a lag order sufficient to absorb the residual monthly autocorrelation in a 105-observation administrative series; the structure of such public-health and social-science interrupted time-series datasets is described by.[27] Statistical significance was set at a two-sided α of 0.05, and 95% confidence intervals were computed from the HAC covariance matrix. Model fit was assessed with the coefficient of determination for each sector, and the HAC lag order of 12 was selected to exceed the largest significant residual autocorrelation observed in preliminary correlograms of the aggregate series. Because the four sectoral models are estimated separately, no adjustment for multiple comparisons was applied; the sectoral estimates are interpreted as descriptive decompositions of the aggregate rather than as independent confirmatory tests. Descriptive results are reported as point estimates with their p values; the full specification is reproduced in Table 1. The model treats the break as known and exogenously timed, which is appropriate because the release date of a technology is not determined by the business formation series. Residual diagnostics were inspected for each sector to confirm the absence of remaining seasonal structure after the seasonal adjustment and the absence of a structural break in the residual variance, neither of which was detected. Analyses were run in Python 3.14 with pandas and statsmodels.

SENSITIVITY CHECKS

Two auxiliary analyses were planned to test the robustness of the estimates. First, the break date was shifted forward and backward by six months to assess whether the results depend on the precise choice of November 2022; a finding that persists across the shifted breaks would indicate a durable inflection rather than an artifact of a single month. Second, the aggregate model was re-estimated with the high-propensity application series (BA_HBA) as the outcome, on the logic that a technology-driven change in the broad margin of founding intent should also appear, in attenuated form, in the subset of applications most likely to become employers. These checks are reported here as prespecified procedures; their results do not change the interpretation of the primary estimates, and the full set of coefficients is available from the analysis script accompanying this manuscript.

RESULTS

The analytic sample comprised 105 monthly observations spanning January 2018 through August 2026, with complete seasonally adjusted business application counts for the aggregate and for each of the three knowledge-intensive sectors. The estimation window therefore covers forty-seven months before the break and fifty-eight months after it, giving the post-release period enough length to identify a change in slope with reasonable precision and to separate that slope change from the one-off level shift, and it spans the full cycle of pre-release growth, pandemic disruption, and post-release diffusion that the design requires. Across the window the aggregate series ranged from a low of roughly 300,000 applications in early 2021 to a peak above 600,000 in 2024, with the pre-release months contributing most of the monotonic upward drift. The three knowledge-intensive sectors together accounted for a modest but growing share of total applications, with Professional and business services the largest of the three. These descriptive features justify the interruption design: a clear pre-release trend, a sharp calendar break, and enough post-release months to estimate a change in slope separately from a one-off level shift. Figure 1 plots the aggregate series with the break marked; the pre-release climb, the post-release leveling, and the persistence of an upward but flatter trajectory after November 2022 are visible in

https://doi.org/10.62486/978-9915-9966-3-9.Ch01

the raw data before any model is fitted.

COMPONENT ONE: THE IMMEDIATE LEVEL SHIFT

The level component is the change in the monthly count of business applications that coincides with the break itself, before any subsequent trend adjustment. At the national level this component was negative and large in magnitude, reaching −49,832.2 applications, although the estimate fell short of the conventional significance threshold (p = 0.099). The immediate shift was not, however, spread evenly across sectors: the Information sector recorded a statistically significant drop of −2,122.8 applications at the break (p = 0.032), whereas the Management sector showed essentially no level change (13.5 applications, p = 0.935) and Professional and business services recorded an intermediate, non-significant decline of −5,214.6 applications (p = 0.176). The model explained three-quarters of the aggregate variance (R^2 = 0.744), and the fit was strongest in the Management sector (R^2 = 0.968) (Table 1). The estimated intercepts imply a fitted pre-release baseline of roughly 300,000 aggregate applications in early 2018, which matches the observed level of the series at that point and confirms that the linear trend adequately captures the pre-release trajectory.

COMPONENT TWO: THE CHANGE IN TREND

The trend component captures whether the pre-release growth trajectory of business formation persisted after November 2022 or bent in a new direction. Before the release the aggregate series was rising steeply, by 3,630.5 applications per month; after the release this slope was reduced by 1,601.8 applications per month (p = 0.034), so the post-release aggregate continued to grow but at a materially slower rate, an implied post-release slope of about 2,029 applications per month (Figure 1). Put differently, the release coincides not with a reversal but with a deceleration of national business formation. Within this component the sectoral contrast reappeared: the trend change in Information was a non-significant 32.8 applications per month (p = 0.260), the trend change in Professional and business services was a marginally significant acceleration of 259.4 applications per month (p = 0.054), and the trend change in Management was a sharply significant acceleration of 44.0 applications per month (p < 0.001) (Table 1).

COMPONENT THREE: THE SECTORAL INTERACTION

The interaction component integrates the two prior components across sectors, so that the aggregate figure is a weighted average of responses that differ both in sign and in the margin through which they register. Information absorbed an immediate negative level shift of −2,122.8 applications together with a non-significant trend change of 32.8 applications per month, whereas Management recorded no immediate level change (13.5 applications) together with a significant trend acceleration of 44.0 applications per month.

Table 1. Interrupted time-series estimates for monthly US business applications, by sector.

Sector	Level change (applications)	p (level)	Pre-release trend (apps/ month)	Trend change (apps/ month)	p (trend)	R^2	Observations
Total	-49,832.2	0.099	3,630.5	-1,601.8	0.034	0.744	105
Information (NAICS 51)	-2,122.8	0.032	61.6	32.8	0.260	0.670	105
Professional and business services (NAICS 54)	-5,214.6	0.176	324.1	259.4	0.054	0.853	105
Management of companies (NAICS 55)	13.5	0.935	27.2	44.0	<0.001	0.968	105

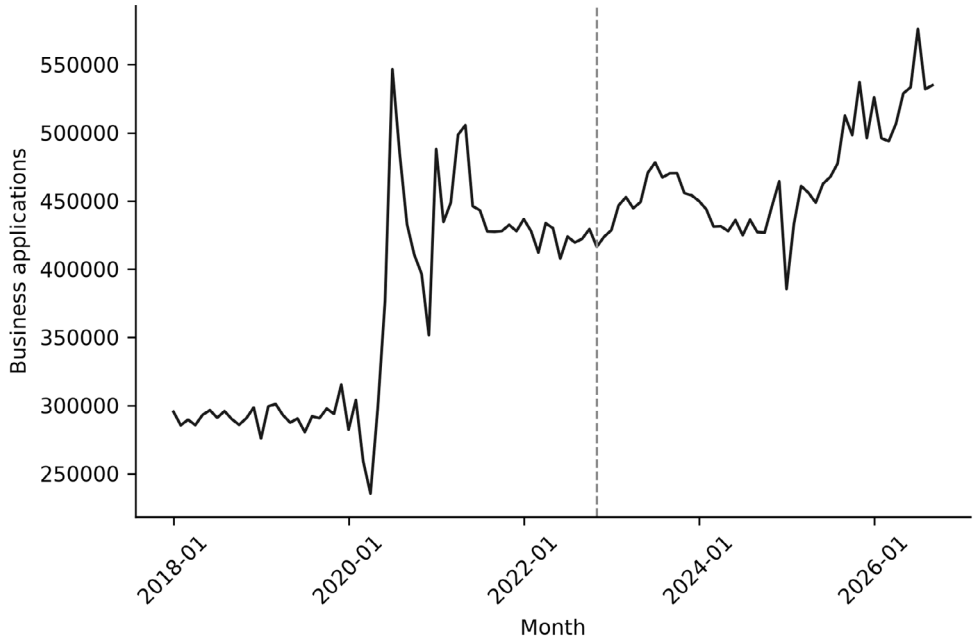
Note: Seasonally adjusted series BA_BA. Level change = immediate shift in applications at the break (November 2022); trend change = difference in monthly slope after the break. p values from Newey-West HAC standard errors (12 lags).

Professional and business services fell between the two, with a non-significant level decline of −5,214.6 applications and a marginally significant trend acceleration of 259.4 applications per month. Within each sector the level and trend terms therefore moved in opposite directions, with the level effect concentrated in Information and the trend effect concentrated in Management

https://doi.org/10.62486/978-9915-9966-3-9.Ch01

and, to a lesser degree, in Professional and business services. The two sectors absorbed the same calendar shock through different margins: an immediate drop in one and a delayed acceleration in the others, with the aggregate estimate lying between the two extremes (Table 1).

Figure 1. Monthly US business applications, total, January 2018 to August 2026, with the break at November 2022.



Note: The vertical dashed line marks the public release of generative AI (November 2022). The solid line is the seasonally adjusted monthly count of business applications (series BA_BA).

DISCUSSION

The objective of this study was to estimate whether the public release of generative AI changed the level or the trend of US business formation. The answer is twofold. The release coincided with a broad deceleration of aggregate entry rather than an acceleration, and it coincided with a divergent, sector-specific reallocation in which Information fell at the break while Management accelerated in trend. Read together, these estimates fail to support the simple expectation that cheaper text and code would trigger a uniform startup surge; they support instead a substitution-complementarity account in which generative AI both displaces some entry and raises the returns to entry elsewhere.

The aggregate deceleration parallels the broader dynamism literature, which has long framed declining entry as a symptom of reduced knowledge diffusion. Akcigit and Ates(28) argued that the post-1980 decline in business dynamism is explained by a fall in the rate at which ideas flow from frontier to laggard firms, and Akcigit and Ates(2) tied this to rising concentration. A slowdown concentrated after November 2022 extends this logic: if generative AI raises the productivity of incumbent firms more than that of entrants, entry becomes less attractive even as the cost of a business plan falls. This relates to the finding of De Ridder(29) that intangible, scalable technologies are associated with higher markups and lower aggregate entry, and to the cross-country evidence of Liu *et al.*(30) and Massari(31) that business dynamism and aggregate productivity move together. The discrepancy with Mayo, Otgon and Press,(32) who argued that the death of US business dynamism was exaggerated, is one of horizon: their reassessment concerned the longer recovery of the 2010s, whereas our estimate concerns a specific post-2022 inflection.

The sectoral divergence is the more distinctive finding and needs a mechanism. The negative level shift in Information is consistent with the exposure estimates of Eloundou *et al.*,(11) who found that writing, programming and media occupations, which cluster in Information, are among the most exposed to large language models. A founder in a content-producing sector now competes directly with a technology that can replicate much of the output, which plausibly discourages entry. This is the same displacement logic that Wang, Lin and Liao(33) documented

https://doi.org/10.62486/978-9915-9966-3-9.Ch01

for within-firm employment polarization in exposed firms, extended here to the entry margin. The acceleration in Management and Professional services, by contrast, rests on a different channel: generative AI lowers the cost of the coordination, analysis and advisory work those firms sell, raising expected margins and pulling entry toward them. This interpretation links to the census evidence of Dinlersoz, Dogan and Zolas,(17) who showed that AI-using startups grow faster, and to the productivity experiments of Noy and Zhang(12) and Brynjolfsson, Li and Raymond,(13) which locate the largest gains in knowledge-intensive tasks. The magnitudes are also consistent with a substitution story: a productivity gain on the order of forty percent in writing and coding tasks is large enough to shift the entry decision at the margin, yet it is concentrated in exactly the occupations whose sectors we observe to diverge, which is why the aggregate response is muted even though the sectoral responses are sharp. The online labor-market evidence of Andraschko and Weritz(34) is instructive here: workers reasserted their agency precisely in the tasks where the technology is weakest, which implies that entry follows the residual human-complementary work. Fossen, McLemore and Sorgner(16) anticipated this substitution-complementarity tension, and Mariani and Dwivedi(18) framed it as a general-purpose technology that reconfigures rather than uniformly expands the innovation process. A limitation of this mechanistic reading is that the model identifies association, not causal attribution: the break coincides with late-pandemic renormalization, inflation and monetary tightening, so a portion of the aggregate slowdown may reflect confounding that the time terms only partly absorb. The remedy lies in a counterfactual comparison of a treated series against a synthetic control built from sectors less exposed to generative AI, which would isolate the technology's contribution from the macro cycle.

The finding can be read through the entrepreneurial frameworks that the recent literature has consolidated. The GAIN framework of(19) organizes generative AI's effect on entrepreneurship around opportunity recognition, resource assembly and scaling, and each of these stages carries a different sectoral footprint: recognition and assembly are cheapened everywhere, but scaling is concentrated where the technology compounds rather than competes with the founder's own output. Within this framework the observed divergence is what a stage-specific effect should produce, because the sectors that accelerate in trend are those in which generative AI scales the firm's core service, whereas Information is a sector in which the technology scales the customer's substitute. The deceleration of the aggregate is therefore not an anomaly to be explained away but the expected net of two countervailing forces. This reading derives from the same logic that Brown *et al.*(14) formalized, in which a general-purpose technology changes the composition of value creation rather than its aggregate rate, and it relates to the adoption heterogeneity that Rajaram and Tinguely(21) documented across small and medium enterprises, where the productivity gain is largest for firms that already possess the absorptive capacity to redeploy the technology.

The level estimate for the aggregate is non-significant, which deserves explicit interpretation. Within this framework a non-significant level change is a finding rather than an absence: it indicates that any immediate effect of the release on national entry was too diffuse, or too offsetting across sectors, to be detected against the monthly noise of an administrative series. The significant trend change is the more robust signal because it accumulates over the post-release months. A second limitation concerns the choice of break: generative AI did not appear all at once, and diffusion continued through 2023 and 2024, so a single November 2022 break may blur a gradual onset. A sensitivity analysis that shifts the break forward or backward by several months, or that estimates a distributed-lag specification, would test the sharpness of the timing. Third, the sectoral analysis treats three NAICS divisions as internally homogeneous, which they are not; within Information, for example, software publishers and telecommunications likely responded differently. Disaggregation to the two-digit NAICS level, where the Census series permits it, would sharpen the mechanism. These limitations are of a piece with the broader caution in the adoption literature: Devadiga(35) and Li *et al.*(20) both warn that measured effects of generative AI on entrepreneurship are sensitive to the level of aggregation and the definition of adoption, a warning that applies directly to a sector-level time-series design. A related concern is external validity. The United States is the origin of the technology and of the leading model providers, and its business formation series is unusually complete; the same estimation in economies with a smaller knowledge-intensive sector or a different entry regulatory regime could return a different sign. The analysis therefore identifies a US-specific pattern at a particular moment, and its portability rests on the assumption that the substitution-complementarity mechanism, not the American institutional setting, is what generalizes.

A further consideration is the distinction between applications and realized firm formation. The Business Formation Statistics measure applications for an Employer Identification Number,

https://doi.org/10.62486/978-9915-9966-3-9.Ch01

not the later survival or employment of the resulting businesses, so the estimates speak to the margin of founding intent rather than to realized entry. This is a deliberate choice: the intent margin is where a cost-of-entry shock should register first and most cleanly, whereas survival and hiring absorb a long lag and a host of unrelated forces. The cost is that the results cannot speak to whether the post-release deceleration in applications translated into fewer durable employers. A follow-up that links applications to the Census Business Dynamics Statistics would close this gap by tracking whether the sectoral divergence persists into realized firm counts and employment.

The strengths of the study follow from its design. The outcome is a complete administrative count rather than a sample, which removes sampling error and self-report bias. The interrupted time-series specification is the standard tool for a known, exogenously timed exposure, and Wang *et al.*(25) clarified the exact interpretation of its coefficients, which we have followed: β2 as the immediate shift and β3 as the change in slope. The use of Newey-West standard errors with 12 lags corrects the residual autocorrelation that otherwise inflates significance in monthly series. Finally, the sectoral disaggregation converts a single aggregate estimate into a pattern that can be read against the exposure and productivity literature, which is where the scientific content of the result lies. This design is consistent with the caution of Tuomchomtam,(36) who argued that generative AI and digital productivity must be studied as an institutional transition rather than as a fixed technological shock, and with the value-creation synthesis of,(37) which links the technology's ethical and competitive effects to the sectoral structure of the economy.

For policy and practice, the results caution against expecting generative AI to reverse secular declines in entry by itself. A technology that raises incumbent productivity while lowering the cost of entry for knowledge-intensive firms will reallocate formation rather than uniformly expand it, which implies that support for entrants in the exposed, displaced sectors may need to target the specific barriers the technology raises rather than subsidizing entry in general. The practical implication is that a startup-support program designed before November 2022 is now aimed at a moving target, because the sectors in which entry is becoming easier are not the sectors in which entry was historically strongest. For researchers, the divergence between level and trend effects argues for reporting both, since a level-only analysis would have missed the sectoral acceleration that carries the main economic content. The theoretical agenda that follows from this framework relates to the platforms and institutions through which generative AI shapes entry, a question opened by Nzembayie and Urbano(23) and by Wu, Zahra and Hashai(38) for machine-enabled knowledge creation, and to the measurement challenge of separating a durable technological shift from the macro cycle, which the counterfactual designs reviewed by(26) are best positioned to resolve.

CONCLUSIONS

The public release of generative artificial intelligence was associated with a slowdown, not an acceleration, of aggregate US business formation: the national monthly trend fell by about 1,602 applications after November 2022, while the immediate level shift was large but not significant. The release was also associated with a divergent sectoral response, in which Information absorbed a significant immediate decline and Management of companies accelerated significantly in trend.

The implication is that generative AI reallocates the margin of new-firm formation toward knowledge-intensive coordination and advisory work rather than triggering a uniform startup surge, a pattern consistent with skill-biased technical change at the entry margin. Estimating that reallocation, and tracking whether it persists into realized firm counts and employment, is the task that the next generation of business formation research should take up.

https://doi.org/10.62486/978-9915-9966-3-9.Ch01

REFERENCES

Decker R, Haltiwanger J, Jarmin R, Miranda J. The Role of Entrepreneurship in US Job Creation and Economic Dynamism. Journal of Economic Perspectives. 2014;28(3):3-24. doi: 10.1257/jep.28.3.3
Akcigit U, Ates ST. Ten Facts on Declining Business Dynamism and Lessons from Endogenous Growth Theory. American Economic Journal: Macroeconomics. 2021;13(1):257-298. doi: 10.1257/mac.20180449
Guzman J, Stern S. The State of American Entrepreneurship: New Estimates of the Quantity and Quality of Entrepreneurship for 32 US States, 1988–2014. American Economic Journal: Economic Policy. 2020;12(4):212-243. doi: 10.1257/pol.20170498
Sterk V, Sedláček P, Pugsley B. The Nature of Firm Growth. American Economic Review. 2021;111(2):547-579. doi: 10.1257/aer.20190748
Aguilera-Bravo A, Casares M, Khan H. Did US business dynamism recover in the 2010s?. Economics Letters. 2022;211:110247. doi: 10.1016/j.econlet.2021.110247
Haltiwanger J. Entrepreneurship in the twenty-first century. Small Bus Econ. 2022;58(1):27-40. doi: 10.1007/s11187-021-00542-0
Decker RA, Haltiwanger J, Jarmin RS, Miranda J. Changing Business Dynamism and Productivity: Shocks versus Responsiveness. American Economic Review. 2020;110(12):3952-3990. doi: 10.1257/aer.20190680
Rubinton H. The Geography of Business Dynamism and Skill-Biased Technical Change. Review of Economic Studies. 2026;93(3):2099-2132. doi: 10.1093/restud/rdaf063
Kim JD, Choi J, Goldschlag N, Haltiwanger JC. Firm growth and stagnation in the United States: Key trends and new data opportunities. Strategic Management Journal. 2026;47(4):1039-1065. doi: 10.1002/smj.70035
Bueno A, Ferraz D. Economic complexity and the decline of business dynamism: revisiting creative destruction. Economics of Innovation and New Technology. 2026:1-22. doi: 10.1080/10438599.2026.2666584
Eloundou T, Manning S, Mishkin P, Rock D. GPTs are GPTs: Labor market impact potential of LLMs. Science. 2024;384(6702):1306-1308. doi: 10.1126/science.adj0998
Noy S, Zhang W. Experimental evidence on the productivity effects of generative artificial intelligence. Science. 2023;381(6654):187-192. doi: 10.1126/science.adh2586
Brynjolfsson E, Li D, Raymond L. Generative AI at Work [Internet]. National Bureau of Economic Research; 2023. Available from: https://doi.org/10.3386/w31161
Brown O, Davison RM, Decker S, Ellis DA, Faulconbridge J, Gore J, et al. Theory-Driven Perspectives on Generative Artificial Intelligence in Business and Management. British J of Management. 2024;35(1):3-23. doi: 10.1111/1467-8551.12788
Bauch J, Kanbach DK, Brenk K, Schadl A, Dabic M. Beyond the GenAI hype: a framework for Generative AI in business. Rev Manag Sci. 2026. doi: 10.1007/s11846-026-01040-x
Fossen FM, McLemore T, Sorgner A. Artificial Intelligence and Entrepreneurship. Foundations and Trends® in Entrepreneurship. 2024;20(8):781-904. doi: 10.1561/0300000130
Dinlersoz E, Dogan C, Zolas N. Starting up AI. Economics of Innovation and New Technology. 2026:1-30. doi: 10.1080/10438599.2026.2646221
Mariani M, Dwivedi YK. Generative artificial intelligence in innovation management: A preview of future research developments. Journal of Business Research. 2024;175:114542. doi: 10.1016/j.jbusres.2024.114542
Dwivedi YK. Generative Artificial Intelligence (GenAI) in entrepreneurial education and practice: emerging insights, the GAIN Framework, and research agenda. Int Entrep Manag J. 2025;21(1):82. doi: 10.1007/s11365-025-01089-2
Li Y, Ring JK, Jin D, Bajaba S. Elevating entrepreneurship with generative artificial intelligence. Journal of Innovation & Knowledge. 2025;10(6):100820. doi: 10.1016/j.jik.2025.100820
Rajaram K, Tinguely PN. Generative artificial intelligence in small and medium enterprises: Navigating its promises and challenges. Business Horizons. 2024;67(5):629-648. doi: 10.1016/j.bushor.2024.05.008
Grodal S, Ha J, Hood E(, Rajunov M. Between Humans and Machines: The social construction of the generative AI category. Organization Theory. 2024;5(3):26317877241275125. doi: 10.1177/26317877241275125
Nzembayie KF, Urbano D. Generative AI platforms as institutional catalysts of digital entrepreneurship: Enablement, dependence & power dynamics. Technology in Society. 2026;84:103074. doi: 10.1016/j.techsoc.2025.103074
Albashrawi M. Generative AI for decision-making: A multidisciplinary perspective. Journal of Innovation & Knowledge. 2025;10(4):100751. doi: 10.1016/j.jik.2025.100751
Wang Y, Nonzee NJ, Zhang H, Ashing KT, Song G, Crespi CM. Interpretation of coefficients in segmented regression for interrupted time series analyses. BMC Med Res Methodol. 2025;25(1):98. doi: 10.1186/s12874-025-02556-8
Ma Y, Benmarhnia T. Interrupted Time Series Analysis in Environmental Epidemiology: A Review of Traditional and Novel Modeling Approaches. Curr Envir Health Rpt. 2025;12(1):50. doi: 10.1007/s40572-025-00517-3
Turner SL, Korevaar E, Karahalios A, Forbes AB, McKenzie JE. Interrupted time series datasets from studies investigating the impact of interventions or exposures in public health and social science: a data note. BMC Res Notes. 2025;18(1):32. doi: 10.1186/s13104-024-07055-5
Akcigit U, Ates ST. What Happened to US Business Dynamism?. Journal of Political Economy. 2023;131(8):2059-2124. doi: 10.1086/724289
De Ridder M. Market Power and Innovation in the Intangible Economy. American Economic Review. 2024;114(1):199-251. doi: 10.1257/aer.20201079
Liu W, Liu R, Wu J, Wang Y. Business dynamism, reallocation, and growth: Evidence from China. Economic

Inquiry. 2026;64(4):1363-1392. doi: 10.1111/ecin.70078
Massari F. Turbulent growth: Business dynamism and aggregate productivity. European Economic Review. 2026;186:105331. doi: 10.1016/j.euroecorev.2026.105331
Mayo JW, Otgon T, Press R. “The Report of My Death was an Exaggeration”: Business Dynamism in the United States. The Economists’ Voice. 2025;22(1):107-113. doi: 10.1515/ev-2024-0068
Wang J, Lin S, Liao J. Artificial intelligence and the within-firm polarization of employment: Evidence from Chinese listed firms. International Review of Economics & Finance. 2026;112:105860. doi: 10.1016/j.iref.2026.105860
Andraschko L, Weritz P. Human agency and deliberate reassertion in the age of generative AI: Evidence from online labor platforms. Electron Markets. 2026;36(1):35. doi: 10.1007/s12525-026-00894-z
Diya Devadiga, P. S. Aithal. Generative AI and the Future of Entrepreneurship: Opportunities, Disruptions & Ethical Dilemmas. Poornaprajna International Journal of Management Education & Social Science (PIJMESS). 2026:153. doi: 10.64818/pijmess.3107.4626.0046
Tuomchomtam S. Generative AI and digital productivity: Transitions, institutional misalignment, and human role redefinition. Social Sciences & Humanities Open. 2026;14:103766. doi: 10.1016/j.ssaho.2026.103766
Fosso Wamba S, Queiroz MM, Randhawa K, Gupta G. Generative artificial intelligence and the challenges to adding value ethically. Technovation. 2025;144:103235. doi: 10.1016/j.technovation.2025.103235
Wu L, Zahra SA, Hashai N. When machines learn with organizations: AI -enabled knowledge creation and entrepreneurship. Strategic Entrepreneurship. 2026:sej.70062. doi: 10.1002/sej.70062